

Abstract of the Disclosure

Disclosed is a data storage architecture employing a plurality of data grids each comprising an array of equal capacity data storage blocks organized into a

5 plurality of rows and a plurality of columns such that each column corresponds to a storage domain in a data storage system and each row of the plurality of rows corresponds to a plurality of data storage blocks, one data storage block from each column of the plurality of columns at the same physical address, with the plurality of rows each having a plurality of contiguous data storage blocks in each domain.

10 Capacity grids are produced by applying one of a plurality of sparing versions that designate at least one data storage block in each row of the grid as designated as spare. Defined within each capacity grid are one or more data storage units. Each data storage unit reflects a storage format, such as a RAID level, and defines data storage blocks, mirror data blocks, and parity data blocks depending upon storage

15 the format and can provide fault tolerance in the event of a domain failure. User data is stored exclusively in a first portion of the grid and parity data, if any, is exclusively stored in a second portion of the grid, providing optimized performance. Data access employs metadata that may comprise grid and storage unit descriptors, data storage unit format, and a sparing table.